

Your GenAI investment: *where the money goes* and how to protect it

Typical project range

\$20K — \$200K

Budget overrun risk

60%

Building a Generative AI solution is one of the highest-leverage investments a company can make right now — and one of the easiest to get wrong. This guide shows you exactly what drives costs, where projects fail, and what separates a high-ROI outcome from a budget spiral.



60% of AI projects built without expert guidance exceed their original budget by 30–50%.

The most expensive project isn't the one with the highest quote — it's the one that starts cheap and balloons through hidden work, scope creep, and rework.

Project Complexity vs. Investment — Where Do You Fit?

1

2

3

4

PoC / Chatbot

API-based · RAG · 2–4 weeks

\$20K–\$30K

Low complexity

Production RAG System

Custom integrations · 1–3 months

\$30K–\$60K

Medium complexity

Fine-Tuned Model

Proprietary data · MLOps · 3–5 mo

\$60K–\$100K

High complexity

Agentic AI Platform

Multi-agent · regulated · 4–8 mo

\$80K–\$200K+

Very high complexity

*Cost of development with an expert partner


What Drives Costs Up — Risk by Approach

APPROACH	EXAMPLE MODELS (2026)	EST. COST WITHOUT EXPERT	MAIN RISK	IN-HOUSE COMPLEXITY
RAG system	Any LLM + vector DB	\$30K–\$70K	Med — data quality	<div style="width: 40%;"></div>
Fine-tune closed-source	GPT-5.4 · Claude Opus 4.6	\$50K–\$130K	Med — vendor lock-in	<div style="width: 45%;"></div>
Fine-tune open-source	Llama 4 · Mistral · DeepSeek R1	\$130K–\$300K+	High — MLOps infra	<div style="width: 60%;"></div>
Agentic AI system	Multi-model · LangGraph · MCP	\$350K+	Very high — safety & governance	<div style="width: 80%;"></div>

*Estimated cost if built without an expert partner — includes hiring, rework, and operational overhead in year one.

Why GenAI is Harder to Budget than Regular Software


Traditional software has predictable costs: engineering hours, hosting, licenses. GenAI introduces entirely new economic variables — mathematical compute, vector storage, GPU cluster rentals, data readiness, and model drift. You're not just paying for code. You're paying for a system that needs to keep working as the world changes.



Data readiness

Your data is almost never ready. Cleaning, labelling, and securing it is the first major cost — and the hardest to estimate upfront.

30–50% of total budget



AI/ML talent shortage

Senior ML engineers command \$160K–\$220K/yr in the US. A full in-house AI team costs \$1M+/yr — before a single model is trained.

Biggest cost for in-house builds



Model selection complexity

Wrong model choice = overpaying for capability you don't need, or underperforming a critical use case. The market changes monthly.

6+ major model families in 2026



Inference costs at scale

100M tokens/month = \$42 on DeepSeek or \$1,000 on GPT-5.4. Traffic spikes can destroy unit economics without smart routing.

Up to 24x cost difference by model



Continuous retraining

Models degrade as the world changes. Without MLOps infrastructure and retraining pipelines, quality drops silently — and expensively.

~\$3K/quarter minimum




Compliance & governance

EU AI Act, GDPR, HIPAA, SOC2. Especially for agentic AI: safety and governance frameworks add \$30K–\$100K before you go live.

Can block entire deployment

6 Mistakes That Turn a \$100K Project Into a \$300K One

These aren't edge cases. They're the patterns that show up on nearly every in-house GenAI project without experienced oversight.



01

PoC code shipped to production

A demo looks great. Leadership approves production. The team inherits PoC code with no error handling, no monitoring, no scalability.


60–80% of budget spent rewriting

02

Scope creep — "just one more feature"

GenAI is flexible, so each expansion sounds small. Cumulatively, a \$60K project becomes \$150K over 6 months of "small additions."

60–150% budget overruns are common



03

Skipping MLOps from day one

Teams ship fast without CI/CD for models, monitoring, or retraining pipelines. Year 2 maintenance costs more than Year 1 development.


2–3x Year 1 costs in Year 2

04

Wrong model for the use case

Using a \$75/M token flagship model for a task that a \$0.30/M model handles fine. Or the reverse — underpowering critical workflows.

Up to 10x unnecessary spend



05

Underestimating data preparation

Teams assume their data is ready. It never is. Messy pipelines, missing labels, and compliance gaps surface mid-project and delay launch.


30–50% of budget, unplanned

06

No ROI measurement from day one

Without baseline metrics defined before build, it's impossible to prove value — or get budget for phase two. Leadership loses confidence.

Project gets shut down at phase 1



The Case for an Expert Partner

The question isn't whether you can build GenAI in-house — it's whether building in-house is the best use of your capital right now. A specialist team brings pattern recognition from dozens of similar projects, pre-built infrastructure, and the experience to avoid the mistakes above. The math usually favors partnership.

VS

Building in-house from scratch

The hidden true cost

- ✗ ML team hiring: 3–6 months + \$160K–\$220K/yr per engineer
- ✗ No existing MLOps infrastructure — build from zero
- ✗ Learning curve: first GenAI project always has costly mistakes
- ✗ No benchmark data — hard to estimate or control budget
- ✗ 60%+ chance of budget overrun on the first project

True year-1 team cost (salaries only)

\$500K–\$1M+

Working with a GenAI specialist

What you actually pay for

- ✓ Ready team: AI architects, ML engineers, MLOps specialists on day 1
- ✓ Pre-built infrastructure, reusable components, proven patterns
- ✓ Experience from 50+ similar projects — known pitfalls avoided
- ✓ Fixed scope gates prevent budget creep
- ✓ 30–40% faster time-to-market vs. in-house hybrid builds

Actual year-1 cost

\$20K–\$200K

35–50%

Cost savings vs. fully in-house team

30–40%

Faster time-to-market with hybrid delivery

3–5x

Higher ROI when project succeeds on first attempt

Get a scoped estimate — before you commit *a single dollar*

Our team will review your use case, data readiness, and compliance requirements — and give you a realistic cost breakdown, risk assessment, and ROI forecast. No vague ranges, no surprises mid-project.

- ✓ Free discovery call — no commitment
- ✓ ROI projection based on your business metrics
- ✓ Data readiness audit included

- ✓ Line-item cost estimate within 5 business days
- ✓ Clear scope gates to prevent budget creep
- ✓ 10+ years of ML & AI delivery experience

Talk to a GenAI Expert →



InData Labs
Big on Data Science & AI

indatalabs.com